
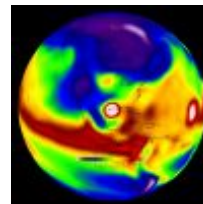

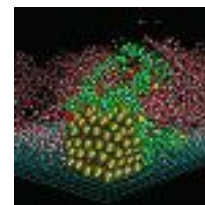


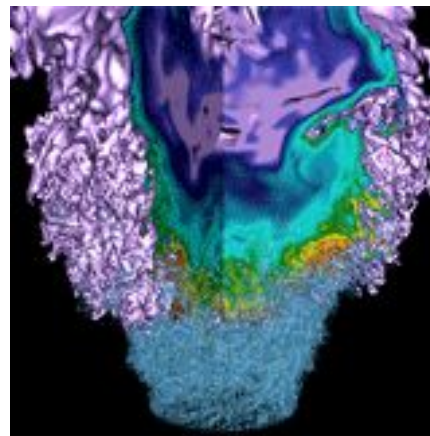


High-Performance Computing and NERSC



Rebecca Hartman-Baker
Group Lead, User Engagement

Brandon Cook
Application Performance


June 10, 2016

High-Performance Computing is






...

... the application of "supercomputers" to scientific computational problems that are either too large for standard computers or would take them too long.

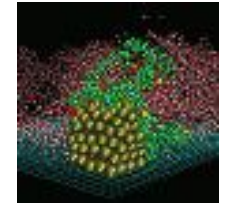

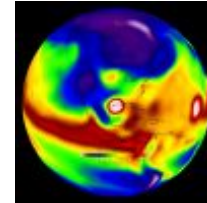


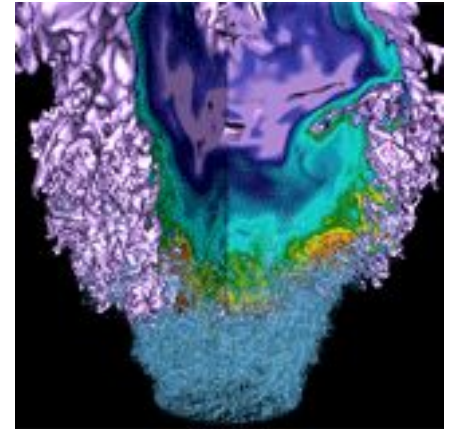


Understanding
How Proteins
Work



Designing Better Batteries

What is a Supercomputer?



What is a supercomputer?




A. A processor (CPU) unimaginably more powerful than the one in my laptop.

B. A quantum computer that takes advantage of the fact that quantum particles can simultaneously exist in a vast number of states.

C. Processors not so different than the one in my laptop, but 100s of thousands of them working together to solve a problem. ✓

A Supercomputer is ...

NERSC



... not so different from a super high-end desktop computer.



Or rather, a lot of super high-end desktop computers.

Edison, shown above, has 5,576 “nodes” (~a powerful desktop), each with 24 compute cores for a total of

133,824 compute cores
 $\sim 2 \times 10^{15}$ calculations/second

But There's More ...

NERSC

The nodes are all connected to each other with a high speed, low latency network.

This is what allows the nodes to “talk” to each other and **work together to solve problems** you could never solve on your laptop or even 150,000 laptops.

Typical point-to-point bandwidth

Supercomputer: 10 GBytes/sec

Your home: 0.02* GBytes/sec

5,000 X

Latency

Supercomputer: 1 μ s

Your home computer: 20,000* μ s

20,000 X



* If you're really lucky



Cloud systems have slower networks

And Even More ...



PBs of fast storage for files and data

Edison: 7.6PB

Your laptop: 0.0005 PB

Your iPhone: 0.00005 PB

15,000 X

Write data to permanent storage


Edison: 140 GB/sec

My iMac: 0.01 GB/sec







14,000 X



Cloud systems have slower I/O and less permanent storage




High-Performance Computing



- implies parallel computing
- In parallel computing, scientists divide a big task into smaller ones
- “Divide and conquer”







For example, to simulate the behavior of Earth's atmosphere, you can divide it into zones and let each processor calculate what happens in each.

From time to time each processor has to send the results of its calculation to its neighbors.



- **This maps well to HPC “distributed memory” systems**
 - Many nodes, each with its own local memory and distinct memory space
 - A node typically has multiple processors, each with multiple compute cores (Edison has 24 cores per node)
 - Nodes communicate over a specialized high-speed, low-latency network
 - SPMD (Single Program Multiple Data) is the most common model
 - Multiple copies of a single program (tasks) execute on different processors, but compute with different data
 - Explicit programming methods (MPI) are used to move data among different tasks

What is NERSC?



National Energy Research Scientific Computing Center (NERSC)



- **NERSC is a national supercomputer center funded by the U.S. Department of Energy Office of Science (SC)**
 - Supports SC research mission
 - Part of Berkeley Lab
- **If you are a researcher with funding from SC, you can use NERSC**
 - Other researchers can apply if research is in SC mission
- **NERSC supports 6,000 users, 800 projects**
- **From 48 states; 65% from universities**
- **Hundreds of users log on each day**


Facility for DOE Office of Science Research




U.S. DEPARTMENT OF
ENERGY

Office of
Science


Largest funder of physical
science research in U.S.




Bio Energy, Environment




Computing




Materials, Chemistry,
Geophysics



Particle Physics,
Astrophysics



Nuclear Physics



Fusion Energy,
Plasma Physics

***NERSC's mission is to accelerate
scientific discovery at the DOE Office of
Science through high performance
computing and data analysis.***

2015 Science Output



2,078 refereed publications



NERSC Nobel Prize Winners



2013 Chemistry



Martin Karplus




2011 Physics



Saul Perlmutter



2006 Physics



George Smoot



2007 Peace



Warren Washington



Nobel Prize in Physics 2015



Scientific Achievement

The discovery that neutrinos have mass and oscillate between different types

Significance and Impact

The discrepancy between predicted and observed solar neutrinos was a mystery for decades. This discovery overturned the Standard Model interpretation of neutrinos as massless particles and resolved the “solar neutrino problem”

Research Details

The Sudbury Neutrino Observatory (SNO) detected all three types (flavors) of neutrinos and showed that when all three were considered, the total flux was in line with predictions. This, together with results from the Super Kamiokande experiment, was proof that neutrinos were oscillating between flavors and therefore had mass



A SNO construction photo shows the spherical vessel that would later be filled with water.

Calculations performed on PDSF & data stored on HPSS played a significant role in the SNO analysis. The SNO team presented an autographed copy of the seminal *Physical Review Letters* article to NERSC staff.

Q. R. Ahmad et al. (SNO Collaboration). Phys. Rev. Lett. 87, 071301 (2001)

Nobel Recipients: Arthur B. McDonald, Queen's University (SNO)
Takaaki Kajita, Tokyo University (Super Kamiokande)

In 2015 scientists at NERSC

used

384,000 single-CPU-years


3,200,000,000

MPP hours of compute time

and currently store

80,000,000

Gbytes of data



Homo erectus
~300,000 years ago

5 million iPhones

Compute Hours



Edison
2,000 M hours



Cori Phase 1
1,000 M hours
Phase 2*
6,000 M hours

* Coming summer 2016

Data Storage



HPSS
75 Petabytes




Project
4 Petabytes

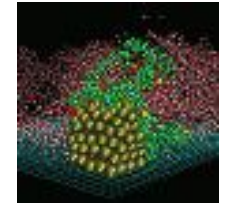

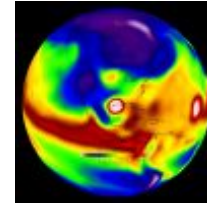


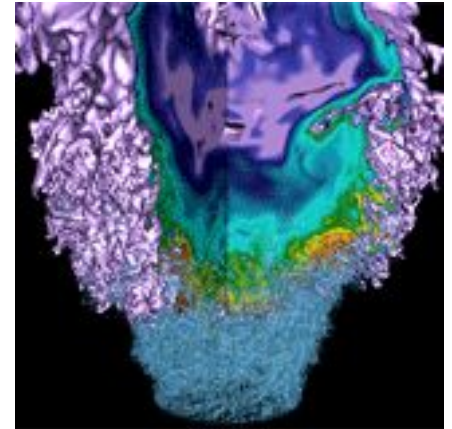
MPP Usage by Science Area



NERSC 2014 Usage by Scientific Discipline





Challenges in HPC



Power: the Biggest Architectural Challenge



If we just kept making computer chips faster and more dense, they'd melt and we couldn't afford or deliver the power.



Now compute cores are getting slower and simpler, but we're getting lots more on a chip.


GPUs and Intel Xeon Phi have 60+ "light-weight cores"

Revolution in Energy Efficiency Needed



Even though energy efficiency is increasing, today's top supercomputer (N=1) uses ~9 MW or roughly \$9M/year to operate. Even if we could build a working exaflop computer today, it would use about 450 MW and cost \$450M/year to pay for power.

Projected Performance Development



Programming for Many-Core: Biggest Software Challenge



- **To effectively use many-core processors, programs must exploit 100K – 1M way parallelism.**
- **Traditional programming paradigms won't work**
 - Too resource intensive per MPI task
 - Data movement is extremely expensive
 - Must incorporate fine-grained (threaded) parallelism
- **Current programming methods for offload to accelerators are difficult and non-portable**
 - Need one “fat core” (at least) for running the OS
 - Data movement from main memory to GPU memory kills performance
 - Programmability is very poor
 - Most codes will require extensive overhauls

Data: Getting Bigger All the Time




Larger simulations are producing ever more data and data subsystems are not keeping up.

Huge experimental facilities like the LHC, beam lines, telescopes, satellites, etc. are producing unprecedented volumes of data at mind-boggling rates.

Reading, writing, and transferring this data is a serious challenge. Making sense of it via data analytics and visualization is too.

Data and job management is another largely unsolved problem. Effective workflows and job schedulers are much needed.



Your Challenges



- **Figure out how to program the next generation of machines**
- **Find a way to make sense of all the data**
- **Build faster, more capable hardware that uses less energy**
- **Create effective data and job management workflows**
- **Bring new fields of science into HPC**
- **Tell the world about what you're doing**



National Energy Research Scientific Computing Center